

In the claims:

1. (original) A method of implementing load balancing in a resilient packet ring (“RPR”) network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise direction and a counterclockwise direction, respectively, wherein adjacent ones of the nodes are connected by two of the links, the method comprising the steps of, for one of the nodes:

determining whether a load imbalance exists at the node in connection with a first class of service; and

responsive to a determination that a load imbalance exists:

changing Bandwidth Broker (“BB”) parameters at the node for the first class of service to cause new flows to be diverted from a more heavily loaded one of the rings to a less heavily loaded one of the rings; and

changing Quality of Service (“QoS”) parameters at the node for the first class of service to improve traffic performance on the more heavily loaded one of the rings, while increasing bandwidth utilization on the less heavily loaded one of the rings.

2. (original) The method of claim 1 wherein the step of determining is performed at periodic time intervals.

3. (original) The method of claim 1 wherein the step of determining is performed using a technique selected from the group consisting of measuring and comparing delays experienced by a test packet sent from the node to a second node via the first and second rings, respectively, and comparing a number of dropped packets on the first and second rings with a preselected maximum value.

4. (original) The method of claim 1 further comprising the step of signaling to a QoS/BB monitor that a load imbalance has been detected responsive to a determination that a load imbalance exists.

5. (original) The method of claim 1 wherein the step of changing the BB parameters comprises the steps of:

decreasing an allocated bandwidth for the first class of service on the more heavily loaded ring; and

increasing an allocated bandwidth for the first class of service on the less heavily loaded ring.

6. (original) The method of claim 1 wherein the step of changing the QoS parameters comprises the steps of:

setting the peak traffic rate to the used bandwidth for the first class of service on the more heavily loaded ring;

reducing token bucket (“TB”) parameters for all other classes of service on the more heavily loaded ring;

setting the peak traffic rate to the used bandwidth for the first class of service on the less heavily loaded ring; and

increasing the number of bytes in a class based queue (“CBQ”) for the first class of service drained off in each scheduler rotation for each of the rings.

7. (original) The method of claim 1 wherein the RPR network is a wavelength division multiplex RPR and the first and second rings are first and second wavelengths, respectively.

8. (original) Apparatus for implementing load balancing in a resilient packet ring (“RPR”) network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise direction and a counterclockwise direction, respectively, wherein adjacent ones of the nodes are connected by two of the links, the apparatus comprising, at one of the nodes:

means for detecting at the node a load imbalance in connection with a first class of service;

means responsive to detection at the node of a load imbalance for changing Bandwidth Broker (“BB”) parameters at the node for the first class of service to cause new flows to be diverted from a more heavily loaded one of the rings to a less heavily loaded one of the rings; and

means responsive to detection at the node of a load imbalance for changing Quality of Service (“QoS”) parameters at the node for the first class of service to improve traffic performance on the more heavily loaded one of the rings, while increasing bandwidth utilization on the less heavily loaded one of the rings.

9. (original) The apparatus of claim 8 wherein the means for detecting performs the detecting at periodic time intervals.

10. (original) The apparatus of claim 8 wherein the means for detecting comprises means for measuring and comparing delays experienced by a test packet sent from the node to a second node via the first and second rings.

11. (original) The apparatus of claim 8 wherein the means for detecting comprises means for comparing a number of dropped packets on the first and second rings with a preselected maximum value.

12. (original) The apparatus of claim 8 further comprising means for signaling to a QoS/BB monitor that a load imbalance has been detected.

13. (original) The apparatus of claim 8 wherein the means for changing the BB parameters comprises:

means for decreasing an allocated bandwidth for the first class of service on the more heavily loaded ring; and

means for increasing an allocated bandwidth for the first class of service on the less heavily loaded ring.

14. (original) The apparatus of claim 8 wherein the means for changing the QoS parameters comprises:

means for setting the peak traffic rate to the used bandwidth for the first class of service on the more heavily loaded ring;

means for reducing token bucket (“TB”) parameters for all other classes of service on the more heavily loaded ring;

means for setting the peak traffic rate to the used bandwidth for the first class of service on the less heavily loaded ring; and

means for increasing the number of bytes in a class based queue (“CBQ”) for the first class of service drained off in each scheduler rotation for each of the rings.

15. (original) The apparatus of claim 8 wherein the RPR network is a wavelength division multiplex RPR and the first and second rings are first and second wavelengths, respectively.

16. (original) Apparatus for implementing load balancing in a resilient packet ring (“RPR”) network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise direction and a counterclockwise direction, respectively, wherein adjacent ones of the nodes are connected by two of the links, the apparatus comprising, at one of the nodes:

a Quality of Service/Bandwidth Broker (“QoS/BB”) monitor responsive to detection at the node of a load imbalance in connection with a first class of service for signaling to a BB to change BB parameters at the node for the first class of service to cause new flows to be diverted from a more heavily loaded one of the rings to a less heavily loaded one of the rings and for changing QoS parameters at the node for the first class of service to improve traffic performance on the more heavily loaded one of the rings, while increasing bandwidth utilization on the less heavily loaded one of the rings.

17. (original) The apparatus of claim 16 wherein detection of a load imbalance is accomplished by measuring and comparing delays experienced by a test packet sent from the node to a second node via the first and second rings.

18. (original) The apparatus of claim 16 wherein detection of a load imbalance is accomplished by comparing a number of dropped packets on the first and second rings with a preselected maximum value.

19. (original) The apparatus of claim 16 wherein the QoS/BB monitor is apprised of a load imbalance via an in-band signaling mechanism.

20. (original) The apparatus of claim 16 wherein the QoS/BB monitor is apprised of a load imbalance via an out-of-band signaling mechanism.

21. (original) The apparatus of claim 16 wherein the QoS/BB monitor changes the BB parameters by:

decreasing an allocated bandwidth for the first class of service on the more heavily loaded ring; and

increasing an allocated bandwidth for the first class of service on the less heavily loaded ring.

22. (original) The apparatus of claim 16 wherein the QoS/BB monitor changes the QoS parameters by:

setting the peak traffic rate to the used bandwidth for the first class of service on the more heavily loaded ring;

reducing token bucket (“TB”) parameters for all other classes of service on the more heavily loaded ring;

setting the peak traffic rate to the used bandwidth for the first class of service on the less heavily loaded ring; and

increasing the number of bytes in a class based queue (“CBQ”) for the first class of service drained off in each scheduler rotation for each of the rings.

23. (original) The apparatus of claim 16 wherein the RPR network is a wavelength division multiplex RPR and the first and second rings are first and second wavelengths, respectively.